arXiv:1611.02480v1 [stat.ME] 8 Nov 2016

Quantile Graphical Models: Bayesian Approaches

Nilabja Guha[*,+], Bani K. Mallick[*]

[*] Department of Statistics, Texas A & M University, College Station, TX 77843, USA.

[+]Institute of Scientific Computation, Department of Mathematics, Texas A & M University, College Station, TX 77843, USA.

Abstract

Graphical models are ubiquitous tools to describe the interdependence between variables measured simultaneously such as large-scale gene or protein expression data. Gaussian graphical models (GGMs) are well-established tools for probabilistic exploration of dependence structures using precision matrices and they are generated under a multivariate normal joint distribution. However, they suffer from several shortcomings since they are based on Gaussian distribution assumptions. In this article, we propose a Bayesian quantile based approach for sparse estimation of graphs. We demonstrate that the resulting graph estimation is robust to outliers and applicable under general distributional assumptions. Furthermore, we develop efficient variational Bayes approximations to scale the methods for large data sets. Our methods are applied to a novel cancer proteomics data dataset where-in multiple proteomic antibodies are simultaneously assessed on tumor samples using reverse-phase protein arrays (RPPA) technology. The measured protein expressions have a wide dynamic range resulting in significant non-Gaussian behaviors. Our analysis reveals several biologically relevant protein-protein connections which are missed by existing approaches.

1 Introduction

Probabilistic graphical models are the basic tools to represent dependence structures among multiple variables. They provide a simple way to visualize the structure of a probabilistic model as well as provide insights into the properties of the model, including conditional independence structures. A graph comprises with vertices (nodes) connected by edges (links or arcs). In a probabilistic graphical model, each vertex represents a random variable (single or vector) and the edges express probabilistic relationship between these variables. The graph defines the way the joint distribution over all the random variables can be decomposed into a product of factors contacting subset of the variables. There are two types of probabilistic graphical models: (1) Undirected graphical models where the edges do not carry the directional information (Schäfer and Strimmer, 2005; Dobra et al., 2004; Yuan and Lin, 2007); (2) The other major class of graphical models is the directed graphical models (DAG) or Bayesian networks where the edges of the graphs have a particular directionality which expresses causal relationships between random variables (Friedman, 2004; Segal et al., 2003; Mallick et al. 2009). In this paper, we focus on the undirected graphical models.

One popular tool of undirected graphical models is Gaussian Graphical Models (GGM) which assume that the stochastic variables follow a multivariate normal distribution with a particular structure of the inverse of the covariance matrix, called the precision or the concentration matrix. This precision matrix of the multivariate normal distribution has the interpretation of the conditional dependence. Compared with the marginal dependence, this conditional dependence can capture the direct link between two variables when all other variables are conditioned on. Furthermore, it is usually assumed that one of the variables can be predicted by those of a small subset of other variables. This assumption leads to sparsity (many zeros) in the precision matrix and reduces the problem to well known

covariance selection problems (Dempster, 1972; Wong et al.,2003). Sparse estimation of precision matrix, thus plays a center role in Gaussian graphical model estimation problem (Friedman et al., 2008).

There has been an intense development of Bayesian graphical model literature over the past decade but mainly in a Gaussian graphical model setup. In a Bayesian setting, this joint modeling is based on hierarchical specifications of the covariance matrix (or precision matrix) using global priors on the space of positive-definite matrices, such as inverse Wishart priors or its equivalents. Lauritzen (1996) introduced an equivalent form as the hyper-inverse Wishart distribution. Although that construction enjoys many advantages, such as computational efficiency due to its conjugate formulation and exact calculation of marginal likelihoods (Scott and Carvalho, 2008), it is sometimes inflexible due to its restrictive form. Unrestricted graphical model determination is challenging unless the search space is restricted to decomposable graphs, where the marginal likelihoods are available up to the overall normalizing constants (Giudici, 1996; Roverato, 2000). The marginal likelihoods are used to calculate the posterior probability of each graph, which gives an exact solution for small examples, but a prohibitively large number of graphs for a moderately large P, the number of nodes/variables. Moreover, extension to a non-decomposable graph is nontrivial and computationally expensive using reversible-jump algorithms (Giudici and Green, 1999; Brooks et al, 2003). Atay-Kayis and Massam (2005) proposed a simple efficient Monte Carlo method. There have been several attempts to shrink the covariance/precision matrix via matrix factorizations for unrestricted search over the space of both decomposable and non-decomposable graphs. Barnard et al. (2000) factorized the covariance matrix in terms of standard deviations and correlations, proposed several shrinkage estimators and discussed suitable priors. Wong et al. (2003) expressed the inverse covariance matrix as a product of the inverse partial variances and the matrix of

partial correlations, then used reversible-jump-based Markov chain Monte Carlo (MCMC) algorithms to identify the zeros among the off-diagonal elements. Liechty et al. (2004) proposed flexible modeling schemes using decompositions of the correlation matrix.

An equivalent formulation of GGM is via neighborhood selection through the conditional mean under normality assumption (Peng et al. 2009). The method is based on the conditional distribution of each variable, conditioning on all other variables. In a GGM framework, this conditional distribution is a normal distribution with the conditional mean function linearly related to the other variables. Furthermore, the conditional independence relationship among variables can be inferred by the variable selection techniques of the regression coefficients of the conditional mean function (Meinshausen and Bühlmann (2006)). More specifically if a specific regression coefficient appeared to be zero, the corresponding variables are conditionally independent. Of course, the joint distribution approach and the conditional approach based on linear regressions are essentially equivalent.

Due to ease of computation and the presence of a nice interpretation, the vast majority of works on graphical model selection have been centered around the multivariate Gaussian distribution. However, the inference of a graph can be significantly affected by deviations from the normality and can lead to a wrong graph. The following example, which we discuss in details in section 5 (Example 1), demonstrates the effect of deviation from normality in a simple case. We assume the following structure for a graph with 15 variables/nodes $X_1, \ldots, X_{15}$ with 200 observations from each variable. Here $X_1, \ldots, X_{10}$ is generated from a Cauchy like heavy tailed distribution and X_i, X_j. $i, j \leq 10$ is connected in the network iff $|i - j| < 2$. We select the underlying network by GGM. The fitted and true graphs are given in Figure 1 where index i denotes i th vertex corresponding to X_i, and it is clear that with deviation from Gaussianity we have a large number of falsely detected edges.

The major problem to extend this concept in a more robust non-Gaussian frame work

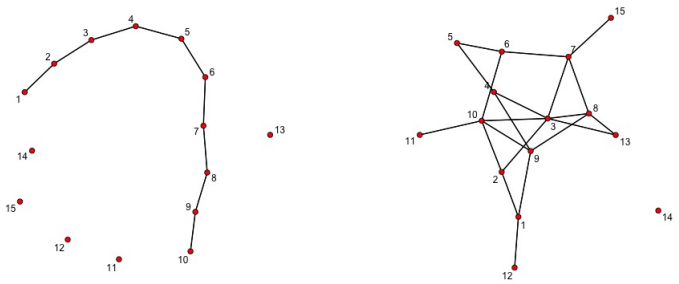

Figure 1: Left panel shows the true network and right hand panel shows the GGM fit for heavy tailed structure, for the example from section 1. Index $1, \ldots, 15$ denotes $X_1, \ldots, X_{15}$.

is that the conditional independence cannot be inferred by the pattern of zero entries in the precision matrix or in the conditional linear regression coefficients, in general. This poses serious restriction in a variety of applications which contain non-Gaussian data as well as data with outliers. Liu et al. (2012) used Gaussian Copula model to allow flexible marginal distributions. Alternatively, non-Gaussian distributions have been directly used for modeling the joint distribution to obtain the graph (Finegold and Drton, 2011; Yang et al., 2014).

In this paper, we propose a novel Bayesian quantile based graphical model. The main intention is to model the conditional quantile functions (rather than the mean) in a regression setup. This is well known that the conditional quantile regression coefficients can infer the conditional independence between variables. Under linearity of the conditional quantile regression function, conditional distribution of the kth variable is independent of the jth variable if the corresponding regression coefficient of the quantile regression is zero for all quantiles. Hence by performing a neighborhood selection of these quantile regression coefficients, we can explore the graphical structure. Thus, in our framework this neighborhood

selection boils down to a variable selection problem in the quantile regression setup. A spike and slab prior formulation has been used for that purpose (George and McCulloch, 1993). The likelihood function depends on a grid of quantiles and borrowing strength from several quantile regression parameters is allowed through a hierarchical Bayesian model. Using Bayesian approach through spike and slab type prior, we can characterize the uncertainty regarding selected graph through the posterior distribution.

A natural development would be to investigate the asymptotic property of the proposed estimated graph. We study the asymptotic behaviors of the graph when the dimension as well as the number of observations increase to infinity. The posterior probability of a small Hellinger neighborhood around the true graph approaches to one, under conditions similar to Jiang (2007). Subsequently, we extend this proof of consistency under the assumption of model misspecification, even under heavy-tailed distribution.

The posterior distribution is not in an explicit form hence we resort to simulation based MCMC method. However, carrying out MCMC in this complex setup could be computationally intensive. Therefore along with MCMC, we also propose a variational algorithm for the mean field approximation of the posterior density (Beal, 2003; Wand et al., 2011; Neville et al., 2014). The main contributions of our paper are: (1) development of robust graphical models based on quantiles in a hierarchical modeling framework, (2) proving the consistency of those resultant graph estimates under truly specified as well as misspecified models and, (3) proposing the MCMC based posterior simulation technique as well as a fast computationally efficient approximation of the posterior distribution.

In next section, we formulate the neighborhood selection problem for a particular node and write down the corresponding likelihood and the posterior density. In section 3, we discuss the estimation consistency. Later in section 4, we discuss the posterior approximation in details and write down the network construction algorithm. In section 5, we

discuss some of the examples and in section 6, we use the proposed method in establishing a protein network.

2 Methodology

An undirected graph G can be represented by the pair (V, E), where V represents the set of *vertices* and $E = (i, j)$ represents the set of edges, for some $i, j \in V$. Two nodes, i and j, are called $n$eighbors if $(i, j) \in E$. A graph is called *complete*, if all possible pair of nodes are $n$eighbors, $(i, j) \in E$ for every $i, j \in V$. $C \subset G$, is called $c$omplete if it induces a complete subgraph. A Gaussian graphical model (GGM) uses a graphical structure to define a set of pairwise conditional independence relationships on a P-dimensional constant mean, normally distributed random vector $\mathbf{x} \sim N_P(\boldsymbol{\mu}, \boldsymbol{\Sigma}_G)$. Here $\boldsymbol{\Sigma}_G$ denotes the dependence of the covariance matrix $\boldsymbol{\Sigma}$ on the graph G and this is the key difference of this class of models with the usual Gaussian models. Thus, if $G = (V, E)$ is an undirected graph and if $\mathbf{x} = (x_\nu)_{\nu \in V}$ is a random vector in $R^{|V|}$ that follows a multivariate normal distribution with mean vector $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}_G$ then the unknown covariance matrix $\boldsymbol{\Sigma}_G$ in GGM is restricted by its Markov properties; given $\boldsymbol{\Omega}_G = {\boldsymbol{\Sigma}_G}^{-1}$, elements x_i and x_j of the vector $\mathbf{x}$ are conditionally independent, given their neighbors, iff $\omega_{ij} = 0$ where w_{ij} is the ijth element of Ω_G. If $G = (V, E)$ is an undirected graph describing the joint distribution of $\mathbf{x}$, $\omega_{ij} = 0$ for all pairs $(i, j) \notin E$. Thus, the elements of the adjacency matrix of the graph G have a very specific interpretation, in the sense that they model conditional independence among the components of the multivariate normal. Presence of an off-diagonal edge in the graph indicates non-zero correlation while its absence indicates zero correlation. This way, the covariance matrix Σ (or the precision matrix Ω) depends on the graph G and this dependence is denoted as Σ_G (Ω_G). The equivalent results can be obtained by using the

conditional regression setup where the conditional distribution of one variable X_k given all other variables is $[X_k|X_{-k}] \sim N(\sum_{j\neq k} \beta_{kj} X_j, \sigma_k^2)$ where $\beta_{kj} = -\omega_{kj}/\omega_{kk}$, $\sigma_k^2 = 1/\omega_{kk}$ and X_{-k} is the vector containing all Xs except the kth one. It is clear that the variable X_k is conditionally independent of X_l given all other variables iff the corresponding conditional regression coefficient β_{kl} is 0. This result transforms the Gaussian graphical model problem to a variable selection (or neighborhood selection) problem in a conditional regression setup (Meinshausen and Bühlmann (2006)).

If multivariate normality assumption on $\mathbf{x}$ does not hold, then the conditional mean does not characterize the dependence among the variables. Under general distribution, it can be helpful to study the full conditional distribution. The absence of an edge between kth and jth node implies that the conditional distribution of X_k given the rest $X_k|X_{-k}$, does not depend on j and vice versa. Any distribution is characterized by its quantiles. Therefore, we can look at the conditional quantile functions of X_k and check if it depends on X_j. Hence, the main idea is to model the quantiles of X_k and perform a variable selection over all quantiles. We use linear model for modeling the quantile functions and perform variable selection in the set up of quantile regression (Koenker and Bassett, 1978; Koenker, 2004) .

Thus, we generalize the concept of Gaussian graphical model in a quantile domain where we consider the conditional quantile regression of each of the node variable X_k given all others say X_{-k} for $k = 1 \cdots P$. In a conditional linear quantile regression model if $X_k(\tau)$ is the τ th quantile of kth variable X_k then the conditional quantile of X_k given X_{-k} can be expressed as

$$X_k(\tau) = \beta_{k,0}(\tau) + \sum_{j\neq k} \beta_{k,j}(\tau) X_j,\ j = 1, \cdots P. \tag{1}$$

We summarize the above discussion in the following result.

Proposition 2.1. *Under the assumption of linearity of the conditional quantile function of X_k, as in model (1), X_k is conditionally independent of X_j iff $\beta_{k,j}(\tau) = 0, \forall \tau$.*

Therefore from Proposition 2.1, we obtain a similar frame work as in the Gaussian graphical model problem. That way, we transform the quantile graphical modeling problem to a quantile regression problem.

Furthermore, instead of looking at a single quantile such as median, considering a set of quantiles will be useful to address a more general dependence structure. To induce sparsity, it will be helpful to look at the coefficients for a set of quantiles τ_s and assume that the condition $\beta_{k,j}(\tau_s) = 0$ for all s implies the conditional independence among the corresponding variables. Indeed, the sparse graphical model based on (1) addresses more general cases than just modeling the conditional mean. In practice instead of the continuum, grid points $0 < \tau_1 < \cdots < \tau_m < 1$ are used for the selection process.

Suppose we have n independent observations which can be presented as a $n \times P$ data matrix $\mathbf{X}^* = \{X_{ij}, i = 1, \cdots, n, j = 1, \cdots, P\}$. We write $\mathbf{X}^* = [X_1, \cdots, X_P]$ where X_i is the $n \times 1$ dimensional ith column vector containing the data corresponding to the ith variable. Since we consider the conditional quantile regression for each of the variable X_j given all the other variables, for the sake of simplicity we describe the general methodology only for a specific variable X_k. For notational convenience, we assume Y is the kth column of $\mathbf{X}^*$ containing the data related to X_k. Furthermore, $\mathbf{X} = \mathbf{X_{-k}}^*$ is a $n \times P$ dimensional matrix containing data corresponding to all other variables except the kth one. Hence, we redefine $\mathbf{X}$ having X_i in the $i + 1$ th column if $i < k$ and X_i in the ith column for $i > k$. We also allow the intercept term as a vector of ones in the first column. In the quantile regression for X_k, we treat Y as the response and $\mathbf{X}$ as the covariates. For the τ th quantile regression, we obtain the estimates of the regression coefficients by minimizing the loss function L such as $\min_{\beta} \sum_{i=1}^{n} \rho_\tau(y_i - \mathbf{x}_i'\beta)$ where $\rho_\tau(z) = z\tau, z \geq 0$ and $\rho_\tau(z) =$

$-(1-\tau)z, z < 0$, the regression coefficient vector $\beta = \{\beta_0, \beta_1, \ldots, \beta_{k-1}, \beta_{k+1}, \ldots, \beta_P\}$, y_i is the ith element of Y and $\mathbf{x}_i$ is the ith row of $\mathbf{X}$.

Mathematically minimizing this loss function L is equivalent to maximizing $-L$ where $\exp(-L)$ is proportional to the likelihood function. This duality between a likelihood and loss, particularly viewing the loss as the negative of the log-likelihood, is referred to in the Bayesian literature as a logarithmic scoring rule (see, for example, Bernardo (1979), page 688). Using loss function to construct likelihood may cause model misspecification. Later we address the issue and show even under model misspecification, we have the posterior concentration around the best linear approximation of the conditional quantile functions. Accordingly, the corresponding likelihood based method can be formulated by developing the model as $y_i = \mathbf{x}_i'\beta + u_i$ where u_is are independent and identically distributed (iid) random variables with the scale parameter t as $f(u|t) = t\tau(1-\tau)exp(-t\rho_\tau(u))$.

Using the results from Li et al.(2009) and Kozumi and Kobayashi (2009), we can express : $u_i = \xi_1 v_i + t^{-\frac{1}{2}}\xi_2\sqrt{v_i}z_i$, where $\xi_1 = \frac{1-2\tau}{\tau(1-\tau)}$, $\xi_2 = \sqrt{\frac{2}{\tau(1-\tau)}}$, $v \sim Exp(t)$ and $z \sim N(0,1)$. Furthermore, the variables indexed by different i s are independent.

The final model can be represented by integrating previous results as

$$\begin{aligned} y_i &= \mathbf{x}_i'\beta + \xi_1 v_i + \xi_2 t^{-\frac{1}{2}}\sqrt{v_i}z_i \\ v_i &\sim Exp(t), \\ z_i &\sim N(0,1). \end{aligned} \tag{2}$$

For selecting the adjacent nodes (neighborhood selection) for node k, a Bayesian variable selection technique has been performed. The stochastic search variable selection (SSVS) is adapted using a spike and slab prior for the regression coefficients as : $p(\beta_j|I_j) =$

$I_j N(0, g^2 v_0^2) + (1 - I_j) N(0, v_0^2)$, (George and McCulloch (1993)) for $j = 1, \ldots, P, j \neq k$ and I_j is the indicator variable related to the inclusion of the jth variable. Let γ be the vector of indicator function I_j's. We denote the spike variance as v_0^2 and the slab variance as $g^2 v_0^2$, where g is a large constant. Alternatively, writing $\beta_{\gamma,j} = \beta_j I_j$ (Kuo and Mallick, 1998) can be helpful, where we use the indicator function in the likelihood and model the quantile of y by $x'\beta_\gamma$. Further, a Beta-Binomial prior is assigned for I_j. The corresponding Bayesian hierarchical model is described as

$$\begin{aligned} \beta_j &\sim N(0, t^{-1}\sigma_\beta^2), \\ I_j &\sim Ber(\pi), \\ \pi &\sim Beta(a_1, b_1), \\ t &\sim Gamma(a_0, b_0). \end{aligned} \tag{3}$$

The Beta Binomial prior opposed to a fixed binomial distribution with a fixed π induces sparse selection (Scott and Berger, 2010).

For a sparse estimation problem we consider m different quantile grid points in $(0, 1)$ as $\tau_1 \ldots, \tau_m$. Let $\underline{\beta}_l = \{\beta_{0,l}, .., \beta_{k-1,l}, \beta_{k+1,l} \ldots \beta_{P,l}\}$ be the coefficient vector corresponding to the τ_l quantile and $\underline{\beta}_{\gamma,l} = \{\beta_{0,l}, .., \beta_{k-1,l} I_{k-1}, \beta_{k+1,l} I_{k+1}, \ldots \beta_{P,l} I_P\}$. Let $\underline{\beta}$ be the vector of all the $\underline{\beta}_l$'s; $\underline{\beta}_{(l-1)P+j} = \beta_{j-1,l}$ if $j < k$ and $\underline{\beta}_{(l-1)P+j} = \beta_{j,l}$ for $j > k$. In this setup, $I_j = 0$ implies that $\beta_{j,l}$ is not in the model and $I_j = 1$ implies that they are included in the model for all l. Thus, we shrink the coefficients for different quantiles as well as induce sparse structure through this group selection. Let t be the scale parameter for τ_l. For τ_l, we write v_i, ξ_1 and ξ_2 from (2) as $v_{i,l}$, $\xi_{1,l}$ and $\xi_{2,l}$, respectively. Let $\mathbf{v}$ be the vector of $v_{i,l}$'s. Using $\tau_1, \tau_2, \ldots, \tau_m$ the corresponding loss function is $l(\underline{\beta}) = \sum_{l=1}^m \rho_{\tau_l}(y_- \mathbf{x}' \underline{\beta}_{\gamma,l})$ and the

corresponding likelihood function is

$$f(y_i|t,\underline{\beta},\gamma) \propto t\exp(-t\sum_{l=1}^{m}\rho_{\tau_l}(y_i-\mathbf{x}_i'\underline{\beta}_{\gamma,l})). \tag{4}$$

The hierarchical model can be written as follows:

$$\begin{aligned}
\underline{\beta}_l|\tau_l &\sim MVN_P(\mathbf{0}_P,\Sigma_{\beta,P\times P}), l=1,\ldots,m,\\
I_j &\sim Ber(\pi),\\
\pi &\sim Beta(a_1,b_1),\\
t &\sim Gamma(a_0,b_0),\\
f(\mathbf{Y}|t,\underline{\beta},\gamma) &\propto t^n\exp(-t\sum_{i=1}^{n}\sum_{l=1}^{m}\rho_{\tau_l}(y_i-\mathbf{x}_i'\underline{\beta}_{\gamma,l})).
\end{aligned} \tag{5}$$

Here, $\mathbf{0}_P$ is a vector of zeros of length P and, $MVN_P(\mathbf{0}_P,\Sigma_{\beta,P\times P})$ denotes P dimensional multivariate normal distribution with the mean vector $\mathbf{0}_P$ and the covariance matrix $\Sigma_{\beta,P\times P}$. We use $\Pi(.)$ to denote prior distributions.

Using the setting in (5) and (2), we can express the posterior distribution of the unknowns as

$$\begin{aligned}
\Pi(\{\underline{\beta}_l,I_j\}_{l,j},\pi,\mathbf{v},t|\mathbf{Y}) \propto t^n\{\prod_{l=1}^{m}t^{\frac{n}{2}}\{\prod_{i=1}^{n}{v_{i,l}}^{-\frac{1}{2}}\exp(-t\frac{(y_i-\mathbf{x}_i'\underline{\beta}_{\gamma,l}-\xi_{1,l}v_{i,l})^2}{2v_{i,l}\xi_{2,l}^2})\times\\
exp(-tv_{i,l})\}\Pi(\underline{\beta}_l)\}\prod_{j\neq k}\Pi(I_j)\Pi(\pi)\Pi(t).
\end{aligned} \tag{6}$$

This posterior distribution is not available in an explicit form and we have to use simulation based approach like Markov Chain Monte Carlo (MCMC) to obtain realizations from it which is described in section 4. Even more, we have to repeat this procedure for

each k, which makes it more computationally demanding. Due to these reasons, we also develop an approximate method based on the variational technique.

3 Graph estimation consistency

In this section, we consider the consistency of the proposed graphical model. Two approaches can be adopted. One method is to look at the variable selection consistency for each of the nodes and the alternative way will be to consider the fitted density induced by the graphical model. We take the latter approach and show the predictive consistency of the proposed network in scenarios encompassing the $P > n$ case. The dimension is adaptively increased with increasing n, the number of observations for each variable. Let $P = p_n$ be the number of nodes. We show that with n increasing to infinity under some appropriate conditions on the prior, the fitted density lies in the Hellinger ball of radius ϵ_n, approaching to zero, around the true density with probability 1, if the proposed model is correct. Next, we consider the case of model miss-specification.

Convergence in exponential rates in terms of Hellinger distance between the posterior graph and the true graph can be achieved under conditions similar to Jiang (2007). Here we briefly define the convergence criterion, describing the conditions required and discuss their implications in terms of the graph estimation.

Let, G^* be the true graph and p_{k,G^*} be the density associated with the k the node of true graph under the proposed model and p_k be the posterior density. Convergence in terms of Hellinger distance such as,

$$P_{G^*}[\pi[d(p_k, p_{k,G^*}) < \epsilon_n | \mathbf{X}^*] > 1 - \delta_n] \geq 1 - \lambda_n$$

where $\epsilon_n, \delta_n, \lambda_n$ going to zero as $n \to \infty$ for each k, can be achieved. Here, $d(p, p^*) = \sqrt{(\int_\chi (\sqrt{p(x)} - \sqrt{p^*(x)})^2 d(\nu(x))}$ denotes the scaled standard Hellinger distance in some measure space χ with measure ν.

For the neighborhood of $Y = X_k$, writing the coefficients $\underline{\beta}_{\gamma,l}^{(-k)} = \underline{\beta}_{\gamma,l}$, $\underline{\beta}_l^{(-k)} = \underline{\beta}_l$, $\beta_{j,l}^{(-k)} = \beta_{j,l}$, $v_{i,l,k} = v_{i,l}$, using (6) we have

$$\Pi(\{\underline{\beta}_l^{(-k)}, I_j^{(-k)}\}_{l,j}, \pi, \mathbf{v}, t|.) \propto \tag{7}$$
$$t^n \{\{\prod_{l=1}^{m} t^{\frac{n}{2}} \{\prod_{i=1}^{n} {v_{i,l,k}}^{-\frac{1}{2}} \exp(-t \frac{(X_{i,k} - \mathbf{x}^*_{-\mathbf{ki}}{}' \underline{\beta}_{\gamma,l}^{(-k)} - \xi_{1,l} v_{i,l,k})^2}{2 v_{i,l} \xi_{2,l}^2}) \times$$
$$exp(-t v_{i,l,k})\} \Pi(\underline{\beta}_l^{(-k)})\} \prod_{j \neq k} \Pi(I_j) \Pi(\pi) \Pi(t)\}$$

where $\mathbf{x}^*_{-ki}$ is the i th row of $\mathbf{X}^*_{-k}$. Through this conditional modeling, we show the posterior concentration of $p(x_k|x_i, i \neq k)p^*(x_i, i \neq k)$ around $p^*(x_1, \ldots, x_P)$.

For the indicator function for the neighborhood selection of k th node, we assume $I_j \sim Ber(\pi_n), j \neq k$, with the restriction $\sum_{j=1}^{p_n} I_j \leq \bar{r}_n$. Let $r_n = p_n \pi_n$. The restriction on the maximum possible dimension can be relaxed by assuming a small probability on the set $\sum_{j=1}^{p_n} I_j \geq \bar{r}_n$. Also, the scale parameter t is assumed to be fixed. The following results also hold for the Beta-Binomial prior on the indicator function and we address it later.

Let ϵ_n be a positive sequence decreasing to zero and $1 \prec n\epsilon_n^2$, where $a_n \prec b_n$ implies $\frac{b_n}{a_n} \to \infty$. Let, $\underline{\beta}^{(-k)}$ be the mP length vector such that $\underline{\beta}_{(l-1)P+j}^{(-k)} = \beta_{j-1,l}^{(-k)}$ if $j < k$ and $\underline{\beta}_{(l-1)P+j}^{(-k)} = \beta_{j,l}^{(-k)}$ otherwise, for $l = 1, \ldots, m$. We have the following prior specifications, $\underline{\beta}^{(-k)} \sim MVN_{mP}(\mathbf{0}_{mP}, S_\beta^{-1})$, where S_β^{-1} is a diagonal matrix in our setting.

Let $\Delta_k^l(r_n) = inf_{|\gamma|=r_n} \sum_l \sum_{j \notin \gamma, j \neq k} |\beta^{*(-k)}_{j,l}|$. Here the superscript '$*$' denotes the true coefficient values. Let $ch_1(M)$ denote the largest eigenvalue of the some positive definite matrix M. Let $\underline{\beta}_\gamma^{(-k)} \sim N(0, V_\gamma)$, i.e the distribution restricted to the variables included in

the model. Let, $B(r_n) = max\{ch_1(V_\gamma), ch_1(V_\gamma^{-1})\}$.

Suppose the following conditions hold.

A1. $\bar{r}_n log(p_n) \prec n\epsilon_n^2$.

A2. $\bar{r}_n log(1/\epsilon_n^2) \prec n\epsilon_n^2$.

A3. $1 \leq r_n \leq \bar{r}_n \leq p_n$.

A4. $\sum_l \sum_{j \neq k} |\beta^{*(-k)}_{j,l}| < \infty$.

A5. $1 \prec r_n \prec p_n < n^\alpha$.

A6. $B(\bar{r}_n) \prec n\epsilon_n^2$.

A7. $p_n \Delta_k^l(r_n) \prec \epsilon_n^2$.

Conditions similar to A1-A7 can be found in Jiang (2007). Condition A1 is needed for establishing the entropy bound on a smaller restricted model space, that is an upper bound on the number of Hellinger balls needed to cover the restricted model space. Conditions A2 and A6 ensure that we have sufficiently large prior probability on the Kullback-Leibler(KL) neighborhood of true model. Assumption A7 is needed to ensure sparsity that is coefficients from all but few variables are close to zero and the total residual effect is small. Also, it has p_n multiplied on the L.H.S as we may not have the boundedness of the node values. Also, the eigenvalue condition is satisfied trivially.

The main idea is to show negligible prior probability for models with dimension larger than $\bar{r}_n$ or where the coefficient vector lies outside a compact set. Then next step would be to cover the smaller model space with $N(\epsilon_n)$ many Hellinger balls of size ϵ_n with $log(N(\epsilon_n)) \prec n\epsilon_n^2$. Tests can be constructed similar to Ghosal et al.(2000). Then by showing that the prior probability of KL neighborhood around the true model has lower bound of some appropriate order, the following results can be achieved.

Let, $h_k = \sqrt{(\int_\chi (\sqrt{p(x_k|x_i, i \neq k)} - \sqrt{p^*(x|x_i, i \neq k)})^2 p^*(x_{i \neq k}) d(\mathbf{x}))}$. Let $d_{t'}(p^*, p) = t'^{-1}\{\int p^*(\frac{p^*}{p})^{t'} - 1\}$ which decreases to KL divergence between p^* and p as t' decreases

to zero. Let the generic term D_n denotes the data matrix. Then we have the following theorem.

Theorem 3.1. *Suppose* $sup_j E|X_j| = M^* < \infty$. *Then from* (6) *under A1-A7, for some* $c_1', c_2' > 0$ *and for* $n^\delta \prec p_n \prec n^\alpha$*;* $\alpha > \delta > 0$ *and for* $sup_{|\gamma| \leq \bar{r}_n}\{ch_1(V_\gamma), ch_1(V_\gamma^{-1})\} \leq B\bar{r}_n^v$*;* $v, B > 0$*, for large enough* $\bar{r}_n$*, the following convergence results hold in terms of the Hellinger distance:*

a)

$$P^*[\Pi(h_k \leq \epsilon_n | D_n) > 1 - e^{-c_1' n\epsilon_n^2}] \to 1.$$

b) Moreover, if we have $\Pi(d_{t'}(p_{k,G^*}, p_k) \leq \epsilon_n) \geq e^{-n\epsilon_n^2}$ *for some* $t' > 0\ \forall k$*, then*

$$P^*[\Pi(h_k \leq \epsilon_n | D_n) > 1 - e^{-c_1' n\epsilon_n^2}; k = 1, \ldots p_n] \geq 1 - n^\alpha e^{-c_2' n\epsilon_n^2}.$$

Proof. Given in the attached supplementary document. □

Remark 3.1. *In particular, for* $\log n \prec r_n \leq \bar{r}_n \prec n^b$ *with* $b = min\{\xi, \delta, \xi/v\}$ *and* $\epsilon_n = n^{-(1-\xi)/2}$ *with* $\xi \in (0, 1)$*, we have* $n\epsilon_n^2 = n^\xi$ *and the convergence rate of the order* e^{-n^ξ}*.*

Remark 3.2. *If each of the node has finitely many neighbors, then some assumptions on tail conditions such as A4, A7 become redundant as only finitely many* $\beta_{j,l}^{*(-k)}$*'s are non zero for each k. For* $p_n = O(n^\alpha), 0 < \alpha < 1$*, we can have* $\bar{r}_n = p_n$ *and* $\epsilon_n = n^{-(1-\xi)/2}$*, with* $\xi \in (\alpha, 1)$*. Thus, we do not need to add any restriction on the model size.*

Remark 3.3. *The results in Theorem 3.1 hold for Beta-Binomial prior on the indicator function as well and given in the attached supplementary document.*

3.1 Consistency under model misspecification

Model (4) has been developed from a loss function and may not be the true data generating model. Therefore, we extend our consistency results under the condition of model misspecification. Let, $f^0_{k,-k}$ be the true density of $Y = X_k$ given $\mathbf{X^*_{-k}}$ and $\mathscr{F}_k$ be the set of densities given by (4). Let $f^*_{k,-k} \in \mathscr{F}_k$ be the density in (4) which has the smallest Kullback-Leibler (KL) distance with $f^0_{k,-k}$. We show that the posterior given in (6) concentrates around $f^*_{k,-k}$. We fix the scale parameter t.

Let $l_{\tau,\underline{\beta}_l^{(-k)}} = E(\rho_{\tau_l}(x_k - x^*_{-k}\underline{\beta}_l^{(-k)}))$ and $\underline{\hat{\beta}}_l^{(-k)} = \arg\min_{\underline{\beta}_l} l_{\tau,\underline{\beta}_l^{(-k)}}$ and suppose the minimizers are unique. Let, $\underline{\hat{\beta}}^{(-k)}$ be the combined vector, analogous to $\underline{\beta}^{(-k)}$. Then under some conditions, the posterior converges to $f_{k,-k}(\underline{\hat{\beta}}^{(-k)})$, the density corresponding to the best linear quantile approximation over $m-$ grid points. Let inf $KL(f^0_{k,-k}, f_{k,-k}(\underline{\beta}^{(-k)})) = \delta^*_k$, which is achieved at the parameter value $\underline{\hat{\beta}}^{(-k)}$.

Posterior concentration under model misspecification needs more involved calculations and can be shown under carefully constructed test functions, as given in Kleijn and Van der Vaart (2006) . However, such approach may depend on the convexity or boundedness of the model space. We take a route similar to Sriram et al. (2013) based on the quantile loss function and show the convergence directly. To prove the consistency, we make a few assumptions. Without loss of generality, we assume that the variables are centered around zero.

Let d_k be the degree (the number of neighbors) of the k th node. Under the following conditions we prove the convergence theorem.

B1. $\max_k d_k < M_0 - 1$ for some universal constant M_0.

B2. $max|X_k| < M > 0$.

B3. There exists $\epsilon > 0$, such that for $|X_k| < \epsilon, \forall k$, $f^{M_0}(X^{M_0}) > \epsilon_1 > 0$ for any M_0 dimensional joint density of X_k's.

B4. $logp_n \prec n$.

B5. $sup_k \|\hat{\underline{\beta}}^{(-k)}\|_\infty < \infty$.

Remark 3.4. *Later we relax condition B2 to include heavy-tailed distributions. If we assume a sub-exponential tail bound on the node variables X_ks and their linear combinations, then the following convergence results hold. We discuss the extension in details in the supplementary document.*

Theorem 3.2. *From* (4) *and* (6)*, under conditions B1–B5, for any $\delta > 0$, $\Pi(KL(f^0_{k,-k}, f_{k,-k}) > \delta + \delta^*_k|.)$ goes to zero almost surely for all k.*

Next, we derive the posterior convergence rate under the model misspecification. For a sequence ϵ_n converging to zero, we assume

B6. $log\frac{1}{\epsilon_n^2} \prec n\epsilon_n^4$,

B7. $\epsilon_n \sim n^{-\xi}, \xi < .25$,

B8. $logp_n \prec n\epsilon_n^4$.

Theorem 3.3. *From* (4) *and* (6)*, under conditions B1–B8, $\Pi(KL(f^0_{k,-k}, f_{k,-k}) > \delta_n + \delta^*_k, \forall k|.)$ goes to zero almost surely, where $\delta_n = 4\epsilon_n^2$.*

Proofs of Theorems 3.2 and 3.3 are given in the supplementary document. We extend our results for heavy-tailed sub-exponential distributions, at the end of the proof of Theorem 3.3.

4 Posterior analysis

We first describe the MCMC steps for posterior simulation. Next, we derive the variational approximation algorithm steps for our case. Let us introduce some notations which will be used in both formulations.

Let $\mathbf{X}_\gamma$ is $n \times P$ dimensional covariate matrix containing $X_i\mathbf{I}_i$ in $i+1$ th column for $i < k$ and $X_i\mathbf{I}_i$ in ith column for $i > k$ and the vector of ones in the first column.

To write the steps for variational approximation and MCMC, we define the following quantities.

- Let $Y_1 = \mathbf{1}_{m\times 1} \otimes Y$ be an $n \times m$ length vector formed by replicating Y, m times.
- Let $\mathbf{X}_{1,\gamma} = I_{m\times m} \otimes \mathbf{X}_\gamma$ and $\mathbf{X}_1 = I_{m\times m} \otimes \mathbf{X}$. Let, $\mathbf{X}^E_{1,\gamma}$ denotes the matrix $\mathbf{X}_{1,\gamma}$, with indicators replaced by their expectations.
- Let $Y^{\delta}_{(l-1)n+i} = Y_{1(l-1)n+i} - \xi_{1,l}(E(\frac{1}{v_{i,l}}))^{-1}$ for $l = 1, \ldots, m$. Similarly, let $Y^{\delta'}_{(l-1)n+i} = Y_{1(l-1)n+i} - \xi_{1,l}(\frac{1}{v_{i,l}})^{-1}$.
- Let Σ be the $nm \times nm$ diagonal matrix, where $(l-1)n+i$ th diagonal entry is $E(t)E(\frac{1}{v_{i,l}\xi^2_{2,l}})$ for $l = 1, \ldots, m$. Similarly, Σ_1 be the $nm \times nm$ diagonal matrix, where $(l-1)n+i$ th diagonal entry is $(t)(\frac{1}{v_{i,l}\xi^2_{2,l}})$ for $l = 1, \ldots, m$.
- Let $S_x = \mathbf{X_1}'\Sigma\mathbf{X_1}$ and $S_{x,\gamma} = \mathbf{X}'_{1,\gamma}\Sigma\mathbf{X}_{1,\gamma}$. Similarly, $S^E_{x,\gamma}$ is the expectation of $S_{x,\gamma}$.
- Let $\underline{\beta}$ be the mP length vector such that $\underline{\beta}_{(l-1)P+j} = \beta_{j-1,l}$ if $j < k$ and $\underline{\beta}_{(l-1)P+j} = \beta_{j,l}$ otherwise, for $l = 1, \ldots, m$. Also, note that we denote the prior for $\underline{\beta}$ as $\underline{\beta} \sim N(0, S^{-1}_\beta)$ as in Section 3.

4.1 MCMC steps

Here, we describe the implementation of the MCMC algorithm to draw realizations from the posterior distribution. More specifically, we use Gibbs sampling by simulating from the complete conditional distributions which are described below (for the k'th node).

(a) For the coefficient vector $\underline{\beta}$:

Given rest of the parameters the conditional distribution is:

$$q^{new}(\underline{\beta}|.) := MVN((S^1_{x,\gamma} + S_\beta)^{-1}(\mathbf{X}_{1,\gamma})'\Sigma_1 Y^{\delta'}, (S^1_{x,\gamma} + S_\beta)^{-1}).$$

Σ_1 be the $nm \times nm$ diagonal matrix, where $(l-1)n + i$ th diagonal entry is $t(\frac{1}{v_{i,l}\xi^2_{2,l}})$ for $l = 1, \ldots, m$ and $S^1_{x,\gamma} = \mathbf{X}_{1,\gamma}\Sigma_1\mathbf{X}_{1,\gamma}$. Note that for the multivariate normal density calculation the matrices are block-diagonal.

(b) For π :

$$q^{new}(\pi|.) := Beta(a_1 + \sum_{j=1, j\neq k}^{P} I_j, P - 1 - \sum_{j=1, j\neq k}^{P} I_j + b_1).$$

(c) For $v_{i,l}$'s:

$$q^{new}(v_{i,j}|.) \propto v_{i,l} f_{InG}(v_{i,l}, \lambda_{i,l}, \mu_{i,l})$$

where $f_{InG}(v_{i,j}, \lambda_{i,l}, \mu_{i,l})$ is a Inverse Gaussian density with parameters $\lambda_{i,l} = t(\frac{(y_i - \mathbf{x}'_i\underline{\beta}_{\gamma,l})^2}{\xi^2_{2,l}})$ and $\mu_{i,l} = \sqrt{\frac{\lambda_{i,l}}{2t + t\frac{\xi^2_{1,l}}{\xi^2_{2,l}}}}$. This step involves a further Metropolis-Hastings sampling with a proposal density for $v_{i,l}$'s as $f_{InG}(v_{i,j}, \lambda_{i,l}, \mu_{i,l})$.

(d) For t:

$$q^{new}(t|.) := Gamma(a_2, b_2)$$

where $a_2 = a_0 + \frac{mn}{2} + n$ and $b_2 = b_0 + \frac{1}{2}\sum_{i,l}(\frac{(y_i - \mathbf{x}_i'\underline{\beta}_{\gamma,l} - \xi_{1,l}v_{i,l})^2}{v_{i,l}\xi_{2,l}^2}) + \sum_{i,l} v_{i,l}$.

(e) For the indicator functions:

$$\log(\frac{P(I_j = 1|.)}{P(I_j = 0|.)}) = (\log\frac{\pi}{1-\pi}) - \frac{1}{2}\{\sum_{i,l,I_j=1} t(\frac{(y_i - \mathbf{x}_i'\underline{\beta}_{\gamma,l} - \xi_{1,l}v_{i,l})^2}{v_{i,l}\xi_{2,l}^2}) -$$

$$\sum_{i,l,I_j=0} t(\frac{(y_i - \mathbf{x}_i'\underline{\beta}_{\gamma,l} - \xi_{1,l}v_{i,l})^2}{v_{i,l}\xi_{2,l}^2})\}.$$

We simulate from this conditional distributions iteratively to obtain the realizations from the joint posterior distribution.

4.2 Variational approximation

As explained in section 2 , we approximate the posterior distribution to facilitate a faster algorithm. We use the variational Bayes methodology for this approximation. First, we briefly review the variational approximation method for posterior estimation. For observed data Y with parameter $\mathbf{\Theta}$ and prior $\Pi(\mathbf{\Theta})$ on it, if we have a joint distribution $p(Y, \Theta)$ and a posterior $\Pi(\mathbf{\Theta}|Y)$ respectively then

$$\begin{aligned}\log p(Y) &= \int \log\frac{p(Y,\Theta)}{\Pi(\mathbf{\Theta}|Y)} q(\mathbf{\Theta}) d(\mathbf{\Theta}) \\ &= \int \log\frac{p(Y,\Theta)}{q(\mathbf{\Theta})} q(\mathbf{\Theta}) d(\mathbf{\Theta}) + KL(q(\mathbf{\Theta}), \Pi(\mathbf{\Theta}|Y)),\end{aligned}$$

for any density $q(\mathbf{\Theta})$. Here $KL(p,q) = E_p(log\frac{p}{q})$, the Kullback-Leibler distance between p and q. Thus,

$$\begin{aligned} \log p(Y) &= KL(q(\mathbf{\Theta}), \Pi(\mathbf{\Theta}|Y)) - KL(q(\mathbf{\Theta}), p(Y, \mathbf{\Theta})) \\ KL(q(\mathbf{\Theta}), p(Y, \mathbf{\Theta})) &= KL(q(\mathbf{\Theta}), \Pi(\mathbf{\Theta}|Y)) - \log p(Y). \end{aligned} \tag{8}$$

With given Y, we minimize $KL(q(\mathbf{\Theta}), p(Y, \mathbf{\Theta}))$. Minimization of the L.H.S of (8) analytically may not be possible in general and therefore, to simplify the problem, it is assumed that the parts of $\mathbf{\Theta}$ are conditionally independent given Y. That is

$$q(\mathbf{\Theta}) = \prod_{i=1}^{s} q_i(\Theta_i)$$

and $\cup_{i=1}^{s} \Theta_i = \mathbf{\Theta}$ is a partition of the set of parameters $\mathbf{\Theta}$. Minimizing under the separability assumption, an approximation of the posterior distribution is computed. Under this assumption of minimizing L.H.S of (8) with respect to $q_i(\Theta_i)$, and keeping the other $q_j(.), j \neq i$ fixed, we develop the following mean field approximation equation:

$$q_i(\Theta_i) \propto exp(E_{-i} \log(p(Y, \mathbf{\Theta})), \tag{9}$$

where E_{-i} denotes the expectation with respect to $q_{-i}(\mathbf{\Theta}) = \prod_{j=1, j\neq i}^{s} q_j(\Theta_j)$. We keep on updating $q_i(.)$'s sequentially until convergence.

For simplicity, we illustrate the variational methodology for $Y = X_{k,n\times 1}$ and $\mathbf{X} = \mathbf{X_{-k}}^*$; i.e the neighborhood selection for the k th node. For notational convenience, we will not use the suffix k in this section and formulate the method for a regression setup. For our case, we have $\mathbf{\Theta} = \{\{\underline{\beta}_l\}, \mathbf{I} = \{I_j\}, \pi, t, \mathbf{v} = \{v_{i,l}\}\}$ with $i = 1, \ldots, n; j = 1, \ldots P, j \neq k, l = 1, \ldots m$.

To proceed, we assume that the posterior distributions of $\{\underline{\beta}_l\}, \mathbf{I} = \{I_j\}, \pi, \mathbf{v} = \{v_{i,l}\}$ and t's are independent given Y. Hence,

$$q(\mathbf{\Theta}) = q(t)q(\underline{\beta})\prod_{j\neq k} q(I_j)\prod_{i,l} q(v_{i,l})q(\pi).$$

Under (3), (6), we have,

$$p(Y,\mathbf{\Theta}) \propto t^n\{\prod_{l=1}^{m} t^{\frac{n}{2}}\{\prod_{i=1}^{n} {v_{i,l}}^{-\frac{1}{2}} \exp(-t\frac{(y_i - \mathbf{x}_i'\underline{\beta}_{\gamma,l} - \xi_{1,l}v_{i,l})^2}{2v_{i,l}\xi_{2,l}^2})exp(-tv_{i,l})\}\}\times$$

$$exp(-\frac{1}{2}(\underline{\beta}'(S_\beta)\underline{\beta}))\pi^{\sum_{j\neq k} I_j + a_1}(1-\pi)^{P-1-\sum_{j\neq k} I_j + b_1} t^{a_0-1}exp(-b_0 t). \quad (10)$$

Using the expression given in (9), we have the variational algorithm given in Table 1.

The densities under this variational approximation algorithm converge very fast and that makes the algorithm many time faster than the standard MCMC algorithms. From (9) we have an explicit form of $q^{new}(.)$ and for our case the updations inside the algorithm are given next.

Sequential updates:

If $q^{old}(\pi), q^{old}(t), q^{old}(v_{i,l}), q^{old}(I_j)$ are the proposed posteriors of $\pi, \{t\}, v_{i,l}$ and I_j's at the current step of iteration, we update

$$q^{new}(\underline{\beta}) \propto exp(E^{old}_{\pi,t,v_{i,l},I_j;i,j,l} log(p(Y,\mathbf{\Theta})))$$

where $E^{old}_{\pi,t,v_{i,l},I_j;i,j,l}$ denotes the expectation with respect to the joint density given by $\prod_{i,j,l} q^{old}(\pi)q^{old}(t)q^{old}(v_{i,l})\ q^{old}(I_j)$.

We have the following closed form expression for updating the densities sequentially.

1. Set the initial values $q^0(\underline{\beta}), q^0(\mathbf{v})$, $q^0(\mathbf{I})$, $q^0(t)$ and $q^0(\pi)$. We denote the current density by $q^{old}()$.
For iteration in 1:N :
2. Find $q^{new}(\underline{\beta})$ by

$$q^{new}(\underline{\beta}) = \arg\min_{q^*(\underline{\beta})} KL\left(q^*(\underline{\beta})q^{old}(\mathbf{I})q^{old}(\mathbf{v})q^{old}(\pi)q^{old}(t), p(Y,\mathbf{\Theta})\right).$$

Initialize $q^{old}(\underline{\beta}) = q^{new}(\underline{\beta})$.
3. Find $q^{new}(\mathbf{I})$ by

$$q^{new}(\mathbf{I}) = \arg\min_{q^*(\mathbf{I})} KL\left(q^{old}(\underline{\beta})q^*(\mathbf{I})q^{old}(\mathbf{v})q^{old}(\pi)q^{old}(t), p(Y,\mathbf{\Theta})\right).$$

Initialize $q^{old}(\mathbf{I}) = q^{new}(\mathbf{I})$.
4. Find $q^{new}(\mathbf{v})$ by

$$q^{new}(\mathbf{v}) = \arg\min_{q^*(\mathbf{v})} KL\left(q^{old}(\underline{\beta})q^{old}(\mathbf{I})q^*(\mathbf{v})q^{old}(\pi)q^{old}(t), p(Y,\mathbf{\Theta})\right).$$

We initialize $q^{old}(\mathbf{v}) = q^{new}(\mathbf{v})$.
5. Find $q^{new}(\pi)$ by

$$q^{new}(\pi) = \arg\min_{q^*(\pi)} KL\left(q^{old}(\underline{\beta})q^{old}(\mathbf{I})q^{old}(\mathbf{v})q^*(\pi)q^{old}(t), p(Y,\mathbf{\Theta})\right).$$

Initialize $q^{old}(\pi) = q^{new}(\pi)$.
6. Find $q^{new}(t)$ by

$$q^{new}(t) = \arg\min_{q^*(t)} KL\left(q^{old}(\underline{\beta})q^{old}(\mathbf{I})q^{old}(\mathbf{v})q^{old}(\pi)q^*(t), p(Y,\mathbf{\Theta})\right).$$

Initialize $q^{old}(t) = q^{new}(t)$.
We continue until the stop criterion is met.
end for
9. Return the approximation $q^{old}(\underline{\beta})q^{old}(\mathbf{I})q^{old}(\mathbf{v})q^{old}(\pi)q^{old}(t)$.

Table 1: Variational Update Algorithm

At each step, the expectations are computed with respect to the current density function.

Thus, for the coefficient vector:

$$q^{new}(\underline{\beta}) := MVN((S^E_{x,\gamma} + S_\beta)^{-1}(\mathbf{X}^E_{1,\gamma})'\Sigma Y^\delta, (S^E_{x,\gamma} + S_\beta)^{-1}).$$

For π we have,

$$q^{new}(\pi) := Beta(a_1 + \sum_{j=1, j\neq k}^{P} E(I_j), P - 1 - \sum_{j=1, j\neq k}^{P} E(I_j) + b_1).$$

For $v_{i,l}$'s

$$q^{new}(v_{i,j}) \propto v_{i,l} f_{InG}(v_{i,l}, \lambda_{i,l}, \mu_{i,l})$$

where $f_{InG}(v_{i,j}, \lambda_{i,l}, \mu_{i,l})$ is a Inverse Gaussian density with parameters $\lambda_{i,l} = E(t)E(\frac{(y_i - \mathbf{x}_i'\underline{\beta}_{\gamma,l})^2}{\xi^2_{2,l}})$ and $\mu_{i,l} = \sqrt{\frac{\lambda_{i,l}}{2E(t) + E(t)\frac{\xi^2_{1,l}}{\xi^2_{2,l}}}}$.

For the tuning parameter t_l,

$$q^{new}(t) := Gamma(a_2, b_2)$$

where $a_2 = a_0 + \frac{mn}{2} + n$ and $b_2 = b_0 + \frac{1}{2}\sum_{i,l} E(\frac{(y_i - \mathbf{x}_i'\underline{\beta}_{\gamma,l} - \xi_{1,l}v_{i,l})^2}{v_{i,l}\xi^2_{2,l}}) + \sum_{i,l} E(v_{i,l})$.

For the indicator function we have

$$\log(\frac{P(I_j = 1)}{P(I_j = 0)}) = E(\log\frac{\pi}{1-\pi}) - \frac{1}{2}\{\sum_{i,l,I_j=1} E(t)E(\frac{(y_i - \mathbf{x}_i'\underline{\beta}_{\gamma,l} - \xi_{1,l}v_{i,l})^2}{v_{i,l}\xi^2_{2,l}}) -$$

$$\sum_{i,l,I_j=0} E(t)E(\frac{(y_i - \mathbf{x}_i'\underline{\beta}_{\gamma,l} - \xi_{1,l}v_{i,l})^2}{v_{i,l}\xi^2_{2,l}})\}.$$

All the moment computations in our algorithm involve standard class of densities. Hence, moments can be explicitly calculated and used in the variational approximation algorithm.

Algorithm for graph construction

Let A be the $P \times P$ adjacency matrix of the target graphical model. Fixing $\tau_1, \ldots, \tau_m$, for $k = 1, \ldots, P$, we compute the posterior neighborhood for each node as follows:

- Construct $Y = X_k$ and $\mathbf{X_{-k}}^*$ as in section 2.
- Compute the posterior of $\Pi(\mathbf{\Theta}|Y)$, by using MCMC or the Variational algorithm, where $\mathbf{\Theta} = \{\{\underline{\beta}_l\}, \mathbf{I} = \{I_j\}, \pi, t, \mathbf{v} = \{v_{i,l}\}\}$ with $i = 1, \ldots, n; j = 1, \ldots P, j \neq k, l = 1, \ldots m$.
- As mentioned earlier in Section 2, $I_j = 0$ implies that $\beta_{j,l}$ is not in the model and $I_j = 1$ implies that they are included in the model for all l. If $P(I_j = 1|Y) > 0.5$ then $A(j,k) = 1$, and $A(j,k) = 0$ otherwise,

Two nodes i and j are connected if at least one of the two is in the neighborhood of the other according to the adjacency matrix A.

5 Some illustrative examples

In this section, we consider two simulation settings to illustrate the application of the proposed methodology. We compare our methodology with the neighborhood selection method for Gaussian graphical model (GGM), by Meinshausen and Bühlmann (2006) (MB), using the R package 'huge'(Zhao et al., 2012) where the model is selected by 'huge.select' function. We also considered graphical Lasso (GLASSO) for graph estimation but in the examples

considered, the MB-neighborhood selection method outperforms GLASSO. Therefore, we consider the MB-neighborhood selection method for our comparison with GGM.

We use $a_1 = 1, b_1 = 1$ in the Beta-Binomial prior. Using the setting of (3) and (10), we use an independent high variance Normal prior on the components of $\underline{\beta}$. The quantile grid points are $\{\tau\} = \{\tau_1, \tau_2, \tau_3\} = \{0.3, 0.5, 0.7\}$ and $\{\tau\} = \{.5\}$. The data is standardized and we choose $t \sim Gamma(1, 1)$.

Example 1

We construct variables $X_1, \ldots, X_{15}$ in the following hierarchical manner using moving average type covariance structure. For $x_{1i}, x_{2i}, \ldots, x_{15,i}$, the i th observation for $X_1, \ldots, X_{15}$, we assume the following hierarchical model: $x_{1,i}, \ldots, x_{10,i} \sim MVN(\mathbf{0}, r_i\Sigma)$, with $\mathbf{0}$ is a vector of zeros and $\Sigma_{kl} = .7^{|k-l|}$ and $\frac{1}{r_i} \sim Gamma(1, 1)$ and $x_{11,i} \ldots, x_{15,i}$ are independent normal variable with mean zero and variance 1. Here $i = 1, \ldots, n$ and variables for different i are independent. Given r_i the k, l th off-diagonal elements of Σ^{-1} are zero if $|k-l| > 2$. Therefore, k and l th node is connected if $|k - l| < 2$. If $r_i = 1, \forall i$, the data is generated from a multivariate normal and the assumption of GGM holds. Choosing r_i in a hierarchical manner creates a heavy tailed structure. We assume that the data is standardized.

Here, $n = 200$ and we use independent normal $N(0, t^{-1}v)$ prior for components of $\underline{\beta}$. We choose $t \sim Gamma(1, 1)$ and $v = 10$. The fit is robust to the selection of v, if we choose larger v ($v \leq 100$). For MCMC we use chain length of 10000 with 4000 burn-ins. A GGM fit for heavy tailed structure along with the true graph is given in Figure 1, as mentioned in section 1. Fitted graphs using different methods are given in Figure 2, where index i denotes the i th vertex corresponding to X_i. It can be seen that using quantile based method we can identify the true graph. Comparing the fitted graph using the neighborhood selection method for Gaussian graphical model, for the heavy tailed misspecified distribution, with

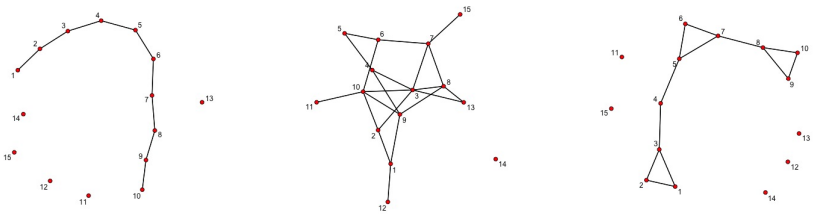

Figure 2: Plots of fitted graphs for example 1: The MCMC fitted network is given in the left panel and the GGM based fitted graph is given in the middle panel. Right most panel shows the GGM fit if the data is generated with $r_i = 1$, i.e under Gaussian assumptions. The variational Bayes approximation coincides with MCMC based fit in this example. Index i denotes the i th vertex corresponding to X_i.

the fit for the case where Gaussian assumption holds ($r_i = 1$), it can be seen that underlying heavy tailed structure has increased the number of falsely detected edges for GGM.

Example 2

To illustrate our method, we consider the following example. We consider $P = 30$ variables $X_1, \ldots, X_{30}$. We construct $X_1, \ldots, X_{10}$ in the following sequential manner:

$$\begin{aligned} X_1, \ldots, X_{10} &\sim Gamma(1, .1) - 10, \\ X_2 &= .4X_1 + \epsilon_1, \\ X_6 &= 1.1X_1 + 4X_4 + 1.3X_9 + \epsilon_2, \\ X_7 &= \Phi^{-1}(\frac{exp(X_2)}{1 + exp(X_2)}) + \epsilon_3, \end{aligned}$$

where $\epsilon_1 \sim N(0,1)$, $\epsilon_2 \sim .5N(2,1) + .5N(-2,1)$, $\epsilon_3 \sim N(0,1)$ and they are independent and independent of the X_i's for each step. The quantity Φ denotes the cdf of standard normal distribution.

Next, we construct $X_{11}, \ldots X_{20}$ from hierarchical multivariate normal random variables $Y_1, \ldots, Y_{10}$, similar to Example 1. That is $y_{1,i}, \ldots, y_{10,i} \sim MVN(\mathbf{0}, r_i\Sigma)$, with $\mathbf{0}$ is a vector of zeros and $\Sigma_{kl} = .7^{|k-l|}$ and $\frac{1}{r_i} \sim Gamma(1,1)$, $i = 1, \ldots, n$, and we have

$$\begin{aligned} X_{11} &= 3Y_3 + 2Y_5 + \epsilon_4, \\ X_{12} &= 3Y_6 + 2Y_7 + \epsilon_5, \\ X_{10+j}, &= Y_j, j \geq 3. \end{aligned}$$

where ϵ_i s are independent and are generated as $\epsilon_4 = .3N(0,9) + .7N(0,1)$, $\epsilon_5 \sim N(0,9)$.

We generate independent normal random variables with mean zero and variance 1 for X_{21} till X_{30}. Hence, we have total number of nodes/variables, $P = 30$ and the graph has two disjoint parts namely: $G_1 = \{X_1, X_2, X_4, X_6, X_7, X_9\}$ and $G_2 = \{X_{11}, \ldots, X_{20}\}$. In addition, non-linear relationships are present between the variables. Generating $n = 200$ independent observations over 100 replications, we construct the network by our algorithm and compare it with the GGM based neighborhood selection method as mentioned earlier. For prior, we use setting similar to Example 1. A typical fit is presented in Figure 3. The quantile based variational Bayes (QVB) method has successfully recovered the connected part inducing sparsity whereas GGM has failed to induce sparseness and estimated the weak connections. Moreover, the quantile based method performs better to separate the independent parts. Using MCMC algorithm, we obtain the similar graphs but the QVB is several hundred time faster than the MCMC. In Table 2, an account of false positivity (detecting an edge, where there is none) has been provided along with the average number

Method	FDR	e_1	e_2	e_3
$QVB(\{\tau\} = \{.5\})$	4.4	1.1	0	0
$QVB(\{\tau\} = \{.3, .5, .7\})$	3	1.3	0.05	0
GGM	11.8	2	0	.25

Table 2: A comparison between GGM and quantile based variational Bayes method(QVB) for example 2.

of undetected edges. Here, FDR denotes the number of falsely detected edges on average per graph, e_1, e_2 and e_3 denote the average number of undetected edges in G_1, G_2 as well as the average number of falsely detected connectors between them. It can be seen that the misspecifications are significantly higher in GGM. The GGM detects a lot of extra edges along with the existing edges. Also, G_1 and G_2 are generally well separated by the quantile based method. In the presence of auto-correlation, some of the edges are not detected by both the methods. Overall, the quantile based variational Bayes provides a sparser and a more accurate solution.

$P > n$ case.

In the next example, we consider a sparse $P > n$ scenario. We consider the previous setup of Example 2 with $n = 100$ and with adding further noise variables $X_{31}, \ldots, X_{120}$ which are generated from a standard normal distribution. Thus, we have $n = 100$ and $P = 120$. The proposed method performs well to detect the true graph as well as provides a sparse solution (see Figure 4). In Figure 4, we also plot the mean residual from the loss function corresponding to the likelihood, with the value of the coefficients as our fitted value from QVB. For the residual plot, we consider $P = 60$ as well, with setting similar to Example 2 and we use 100 replications.

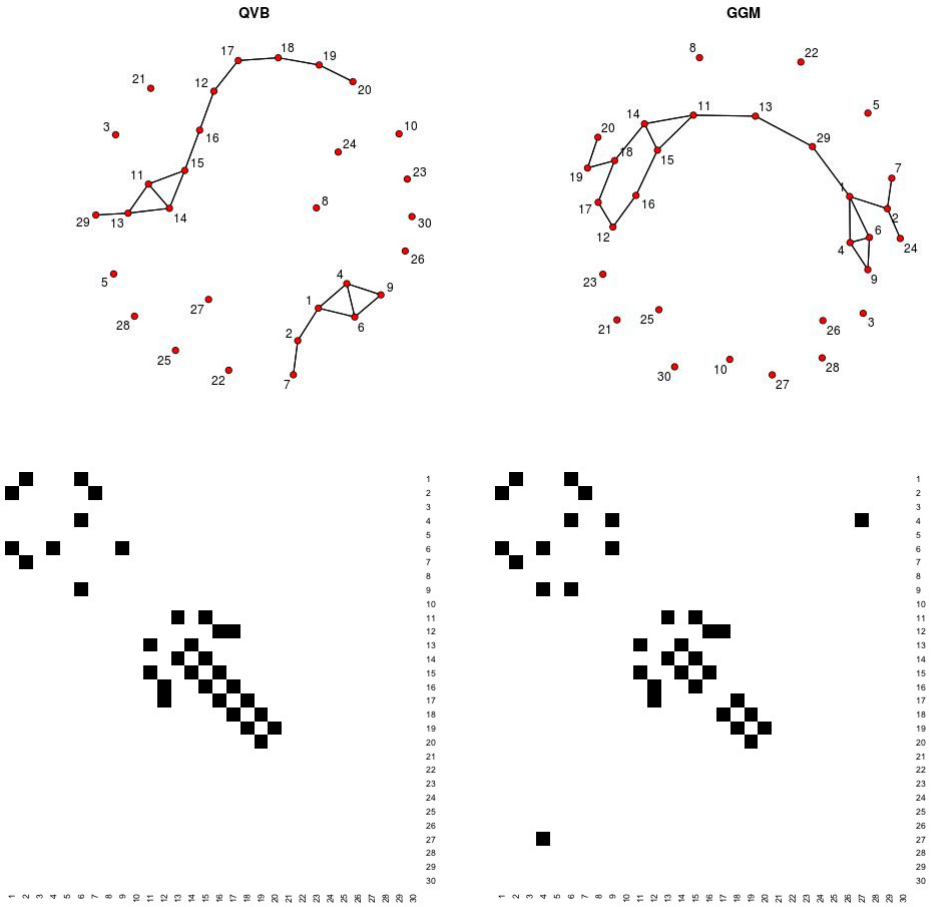

Figure 3: At top panel, left panel side shows the network constructed by the quantile based variational Bayes (QVB) with $\{\tau\} = \{.5\}$ and the right hand panel shows a construction by GGM based method. Bottom panel shows the adjacency matrix of a fitted network for $\tau = \{.3, .5, .7\}$ in the right panel and the true adjacency matrix for the true graph in the left panel. Here black at the i, j place denotes the presence of an edge between i and j th vertices. Index i denotes i th vertex corresponding to X_i.

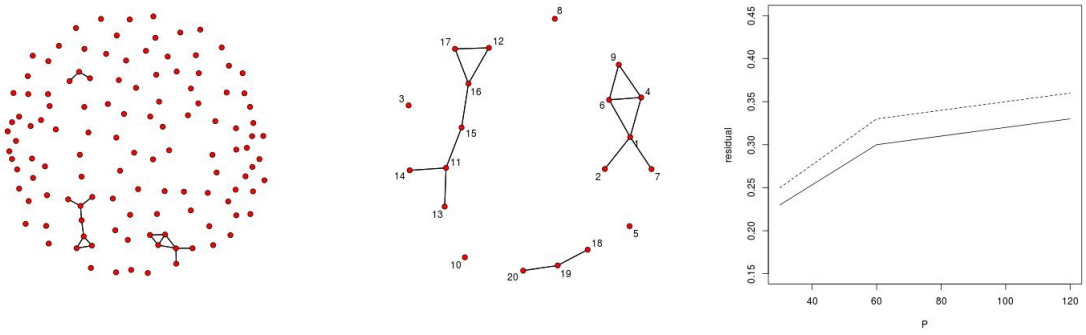

Figure 4: Left panel shows a sparse fitted graph by MCMC for example 2 $P > n$ case ($P = 120, n = 100$), with $\{\tau\} = \{.3, .5, .7\}$ and middle panel shows the connected variables. Right most panel shows the residual plot for different $\{\tau\}$ and P. Here $\{\tau\} = \{.3, .5, .7\}$ is represented by solid line and $\{\tau\} = \{.5\}$ is represented by dotted line.

$\{\tau\}$	$P = 30$	$P = 60$	$P = 100$
		C_E	
$\{\tau\} = \{.4, .6\}$	125	137	160
$\{\tau\} = \{.3, .5, .7\}$	165	177	206

Table 3: The relative computational efficiencies (C_E), the ratio of the runtime of MCMC and of the variational methods are given in the setting of example 1 and 2, with $P = 30, 60$ and 100, and $n = 100$.

Computational gain due to QVB

In all the cases the variational approximation as well as the MCMC based algorithm perform well to detect the true graphs. However, QVB is many times faster than the MCMC. Using 5000 samples for each node, the MCMC runtime is usually more than 100 times of that of the QVB. With higher dimension, the computational gain increases even more. Let C_E be the ratio of the runtime of MCMC and QVB methods. Estimates of C_E are given in Table 3, where the number of quantile grids are 2 and 3.

6 Protein network

The Cancer Genome Atlas (TCGA) is a source of molecular profiles for many different tumor types. Functional protein analysis by reverse-phase protein arrays (RPPA) is included in TCGA and looking at the proteomic characterization the signaling network can be established.

Proteomic data generated by RPPA across 3,467 patient tumors obtained from TCGA includes 11 cancer types. We consider lung squamous cell carcinoma (LUSC) data set. The data set considered, has $n = 121$ observations with $P = 174$ high-quality antibodies. The antibodies encompass major functional and signaling pathways relevant to human cancer and a relevant network gives us their interconnection subject to LUSC. A comprehensive analysis can be found in Akbani et al. (2014) where the EGFR family along with MAPK and MEK lineage was found to be dominant determinant of signaling, where for LUSC it was mainly EGFR.

We use our quantile based approach with $\{\tau_l\} = \{.2, .4, .6, .8\}$ and a flat prior on $\underline{\beta}$ (independent $N(0, 100)$). The graphical LASSO method cannot select a sparse (using huge.select) network. Choosing the penalization by direct cross validation in GLASSO in Akbani et al. (2014), the network has been generated. Here for the LUSC dataset the networks can be compared for new connections and the detection of the known connections. The network and the connection tables with variable index can be found in Figure 5 and Table 4. The type of the connection (positive/negative) is also provided. We can say that one variable effects other variable positively (negatively), conclusively, if the coefficients in the corresponding quantile regression is greater (less) than zero for all the quantiles.

In our fitted network, the strong EGFR/HER2 connections are detected as seen in Akbani et al.. The connection between $EGFRpY1068$ and $HER2p\ Y1248$ is detected

which are known to cross react. The negative connection between $p85$ and $claudin7$ is detected reaffirming its finding in Akbani et al.. The connection between $E.Cadherin$ and $alph/beta.Catenin$ is detected as expected. Similar to Akbani et al., $pAkt$ and $pPras40$ are found to be connected. Also, MEK is active and connected to $PMAPK$.

We have detected some new connections such as between $SNAI2$ and $PARP1$. Here, SNAI2 is a DNA-transcriptional repressor and $PRAP1$ modulates transcription. The connection between $PRAS40$ and $ps473$ is not established for LUSC before. $AKT1S1.PRAS40$ regulates the metabolism and the cell growth responding to nutrients and hormonal signal, whereas $AKTps473$ regulates processes such as metabolism, proliferation, cell survival, growth and angiogenesis.

Proteins, $casp8$ and $ERCC1$ are found to be connected with MET, which are also new connections. $CASP8$ performs protein metabolism and $ERCC1$ is related to structure specific DNA repairing. They both are connected to growth factor receptor MET. The connection between $KRAS$ and $smad4$ is detected, while the connection between $YWHAE$ and $KRAS$ is new and needs further study. All the connections with the corresponding proteins can be found in table 5.

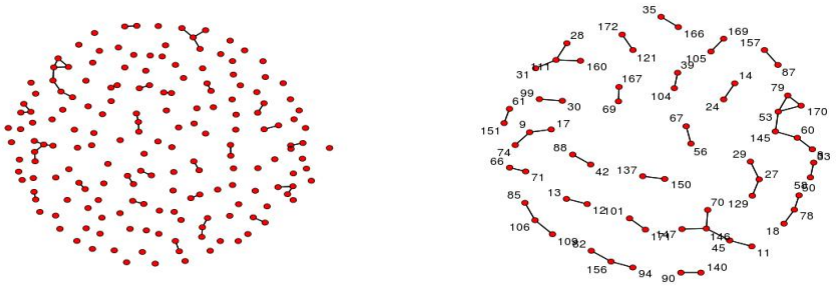

Figure 5: Left hand side shows the network constructed by the quantile based variational algorithm on the LUSC data set and the right hand panel shows the connected part.

Node $i \leftrightarrow$ Node j	Proteins	Sign of connection
$13 \leftrightarrow 12$	$MSH6 \leftrightarrow MSH2$	+
$17 \leftrightarrow 9$	$AktPT308 \leftrightarrow AktPS473$	+
$24 \leftrightarrow 14$	$ACC1 \leftrightarrow ACC$	+
$29 \leftrightarrow 27$	$beta.cantenin \leftrightarrow E.Cadherin$	+
$45 \leftrightarrow 11$	$SNAI2 \leftrightarrow PARP1$	+
$58 \leftrightarrow 8$	$CCND1.Cyclin \leftrightarrow CD20$	+
$60 \leftrightarrow 33$	$SrcpY416 \leftrightarrow Src$	+
$67 \leftrightarrow 56$	$GSK3.alpha.beta.pS21 \leftrightarrow GSk3pS9$	+
$71 \leftrightarrow 66$	$Pai1 \leftrightarrow Fibrocentin$	+
$74 \leftrightarrow 9$	$PRAS40_pT246 \leftrightarrow Akt_pS473$	+
$78 \leftrightarrow 18$	$Chk2_pT68 \leftrightarrow CDKN1_ppT157$	+
$78 \leftrightarrow 50$	$Chk2_pT68 \leftrightarrow chk2.M$	+
$88 \leftrightarrow 42$	$mTORps2448 \leftrightarrow mTOR$	+
$99 \leftrightarrow 30$	$Beclin \leftrightarrow XBP1$	+
$104 \leftrightarrow 39$	$YAP_PS127 \leftrightarrow YAP$	+
$106 \leftrightarrow 85$	$AMPK_alpha \leftrightarrow CDK1$	+
$109 \leftrightarrow 106$	$AMPK_PT172 \leftrightarrow AMPK_alph$	+
$111 \leftrightarrow 28$	$EIF4EBP1_pT37 \leftrightarrow EIF4EBP1$	+
$111 \leftrightarrow 31$	$EIF4EBP1_pT37 \leftrightarrow EIF4EBP1_PS65$	+
$129 \leftrightarrow 27$	$alpha.Catenin \leftrightarrow E.Cadherin$	+
$140 \leftrightarrow 90$	$cKit \leftrightarrow GAB2$	+

Node $i \leftrightarrow$ Node j	Proteins	Sign of connection
$145 \leftrightarrow 53$	$HER2_pY1248 \leftrightarrow EGFR_PY1068$	+
$145 \leftrightarrow 60$	$HER2_pY1248 \leftrightarrow Src_PY416$	+
$146 \leftrightarrow 45$	$MET \leftrightarrow SNAI2$	+
$146 \leftrightarrow 70$	$MET \leftrightarrow CASP8$	+
$147 \leftrightarrow 146$	$ERCC1 \leftrightarrow MET$	+
$150 \leftrightarrow 137$	$PI3k.p85 \leftrightarrow Claudin.7$	-
$151 \leftrightarrow 61$	$BCL2 \leftrightarrow Bim$	+
$156 \leftrightarrow 82$	$KRAS \leftrightarrow YWHAE$	+
$156 \leftrightarrow 94$	$KRAS \leftrightarrow Smad4$	+
$157 \leftrightarrow 87$	$ARID1A \leftrightarrow NCOA3.AIB1$	+
$160 \leftrightarrow 111$	$EIF4EBP1 \leftrightarrow EIF4EBP1_pT37$	+
$166 \leftrightarrow 35$	$EIF4E \leftrightarrow MEK1$	+
$167 \leftrightarrow 69$	$MAPK_pT202_Y204 \leftrightarrow MEK1_pS217_S221$	+
$169 \leftrightarrow 105$	$PKC.alpha_pS657 \leftrightarrow PKC.alpha$	+
$170 \leftrightarrow 53$	$EGFR \leftrightarrow EGFR_pY1068$	+
$170 \leftrightarrow 79$	$EGFR \leftrightarrow EGFR_pY1173$	+
$171 \leftrightarrow 101$	$Rab25 \leftrightarrow SETD2$	+
$172 \leftrightarrow 121$	$Rb_pS807_S811 \leftrightarrow Rb$	+

Table 4: Connections and corresponding nodes in LUSC data set

7 Discussion

The proposed approach offers a robust, non-Gaussian model as well as easily implementable algorithms for sparse graphical models. Even with large values of P and relatively smaller value of n, it is possible to detect underlying connections where other existing methods may fail as shown in example 2 and in the analysis of the LUSC data set. In the protein network construction, we are able to establish the known signaling network with some newly discovered connections, which will be validated in our Bioinformatics Laboratory.

In this development, we prove the model estimation consistency and posterior concentration rate under both the true model and the misspecified model. Under misspecified model the posterior concentration occurs around the minimum KL distance point from the true density and the set of proposed densities. From simulation examples where we do not assume any density structure in the data generating model, the proposed method performs well. In future, we will further investigate the model selection properties for each node and validate if the true model's posterior probability converges to one.

8 References

1. Akbani, R., Ng, P. K. S., Werner, H. M., Shahmoradgoli, M., Zhang, F., Ju, Z., ... , & Ling, S. (2014). A pan-cancer proteomic perspective on The Cancer Genome Atlas. Nature communications, 5.

2. Atay-Kayis, A., & Massam, H. (2005). A Monte-Carlo Method for Computing the Marginal Likelihood in Nondecomposable Gaussian Graphical Models. Biometrika, 92:317–335.

3. Barnard, J., McCulloch, R., & Meng, X. L. (2000). Modeling covariance matrices in

terms of standard deviations and correlations, with application to shrinkage. Statistica Sinica, 10(4), 1281-1312.

4. Beal, M. J. (2003). Variational algorithms for approximate Bayesian inference. Ph.D. thesis, Gatsby Computational Neuroscience Unit, University College London.

5. Bernardo, J. M. (1979). Expected information as expected utility. The Annals of Statistics, 686-690.

6. Brooks, S.P., Giudici, P., & Roberts, G.O.(2003). Efficient construction of reversible jump Markov chain Monte Carlo proposal distributions. J. R. Stat. Soc., Ser. B, Stat. Methodol. 65(1), 3–39.

7. Dempster, A.P. (1972). Covariance Selection. Biometrics 28: 157–175.

8. Diaconis, P., & Ylvisaker, D. (1979). Conjugate Priors for Exponential Families. Annals of Statistics, 7: 269–281.

9. Dobra, A., Hans, C., Jones, B., Nevins, J. R., Yao, G., & West, M. (2004). Sparse graphical models for exploring gene expression data. Journal of Multivariate Analysis, 90(1), 196-212.

10. Finegold, M., & Drton, M. (2011). Robust graphical modeling of gene networks using classical and alternative t-distributions. The Annals of Applied Statistics, 1057-1080.

11. Friedman, N. (2004). Inferring cellular networks using probabilistic graphical models. Science, 303, 799–805.

12. Friedman, J., Hastie, T., & Tibshirani, R. (2008). Sparse inverse covariance estimation with the graphical lasso. Biostatistics 9: 432–441.

13. George, E. I., & McCulloch, R. E. (1993). Variable selection via Gibbs sampling. Journal of the American Statistical Association, 88(423), 881–889.

14. Ghosal, S., Ghosh, J. K., & Van Der Vaart, A. W. (2000). Convergence rates of posterior distributions. Annals of Statistics, 28(2), 500–531.

15. Giudici, P. (1996). Learning in graphical Gaussian models. Bayesian Statistics, 5, 621-628.

16. Giudici, P., & Green, A. P. (1999). Decomposable graphical Gaussian model determination. Biometrika, 86(4), 785-801.

17. Jiang, W. (2005). Bayesian variable selection for high dimensional generalized linear models. Technical Report 05-02, Dept. Statistics, Northwestern Univ.

18. Jiang, W. (2007). Bayesian variable selection for high dimensional generalized linear models: convergence rates of the fitted densities. The Annals of Statistics, 35(4), 1487-1511.

19. Kleijn, B. J., & van der Vaart, A. W. (2006). Misspecification in infinite-dimensional Bayesian statistics. The Annals of Statistics, 837-877.

20. Koenker, R. (2004). Quantile Regression for Longitudinal Data. Journal of Multivariate Analysis, 91: 74–89.

21. Koenker, R., & Bassett Jr, G. (1978). Regression quantiles. Econometrica: journal of the Econometric Society, 33-50.

22. Kozumi, H., & Kobayashi, G. (2011). Gibbs sampling methods for Bayesian quantile regression. Journal of statistical computation and simulation, 81(11), 1565-1578.

23. Kuo, L., & Mallick, B. (1998). Variable selection for regression models. Sankhya Ser. B, 60: 65–81.

24. Lauritzen, S. L. (1996). Graphical Models. Claredon, Oxford.

25. Li,H.,& Gui, J.(2006). Gradient directed regularization for sparse Gaussian concentration graphs with applications to inference of genetic networks. Biostatistics. Vol 7, 302–317.

26. Li, Q., Xi, R., & Lin, N. (2010). Bayesian regularized quantile regression. Bayesian Analysis , 5 , 3: 533–556. doi:10.1214/10-BA521.

27. Liechty, J. C., Liechty, M. W., & Müller, P. (2004). Bayesian correlation estimation. Biometrika, 91(1), 1-14.

28. Liu, H., Han, F., Yuan, M., Lafferty, J.,& Wasserman, L. (2012). High-dimensional semiparametric Gaussian copula graphical models. The Annals of Statistics, 40(4), 2293-2326.

29. Mallick, B., Gold, D., & Baladandayuthapani, V. (2009). Bayesian analysis of Gene expression data. Wiley.

30. Meinshausen, N., & Bühlmann, P. (2006). High-dimensional graphs and variable selection with the lasso. The Annals of Statistics, 1436–1462.

31. Neville, S.E., Ormerod, J.T., & Wand, M.P. (2014) Mean field variational Bayes for continuous sparse signal shrinkage: pitfalls and remedies. Electronic Journal of Statistics, 8, 1113-1151.

32. Peng, J., Wang, P., Zhou, N., & Zhu, J. (2009). Partial correlation estimation by joint sparse regression models. Journal of the American Statistical Association, 104(486).

33. Roverato, A. (2000). Cholesky decomposition of a hyper inverse Wishart matrix. Biometrika, 87(1), 99-112.

34. Schäfer, J., & Strimmer, K. (2005). A shrinkage approach to large-scale covariance matrix estimation and implications for functional genomics. Statistical applications in genetics and molecular biology, 4(1).

35. Scott, J. G., & Carvalho, C. M. (2008). Feature-inclusion stochastic search for Gaussian graphical models. Journal of Computational and Graphical Statistics, 17(4).

36. Scott, J. G., & Berger, J. O. (2010). Bayes and empirical-Bayes multiplicity adjustment in the variable-selection problem. The Annals of Statistics, 38(5), 2587-2619.

37. Segal, E., Shapira, M., Regev, A., Pe'er, D., Botstein, D., Koller, D., & Friedman, N. (2003). Module networks: identifying regulatory modules and their condition-specific regulators from gene expression data. Nature genetics, 34(2), 166-176.

38. Sriram, K., Ramamoorthi, R. V., & Ghosh, P. (2013). Posterior consistency of Bayesian quantile regression based on the misspecified asymmetric Laplace density. Bayesian analysis, 8(2), 479–504.

39. Wand, M.p., Ormerod, J.T., Padoan, S.A., & Fruhwirth, R. (2011). Mean field variational Bayes for elaborate distributions. Bayesian Analysis, 6, 847-900.

40. Wang, H. (2012). Bayesian Graphical Lasso Models and Efficient Posterior Computation. Bayesian Analysis, 7, 771–790.

41. Wong, F., Carter, C. K., & Kohn, R. (2003). Efficient estimation of covariance selection models. Biometrika, 90(4), 809-830.

42. Yang, E., Allen, G., Liu, Z., & Ravikumar, P. K. (2012). Graphical models via generalized linear models. In Advances in Neural Information Processing Systems, 1358-1366.

43. Yuan, M., & Lin, Y. (2007). Model selection and estimation in the Gaussian graphical model. Biometrika, 94(1), 19-35.

44. Zhao, T., Liu, H., Roeder, K., Lafferty, J., & Wasserman, L. (2012). The huge package for high-dimensional undirected graph estimation in R. The Journal of Machine Learning Research, 13(1), 1059–1062.

Supplementary materials

Proofs of Theorem 3.1, 3.2 and 3.3 are given in Section 1 of the attached supplementary document. We also derive the variational updates in Section 2.